

AFRL-IF-RS-TR-2004-224
Final Technical Report
July 2004



RESEARCH NETWORKS AND TECHNOLOGY MIGRATION (RESNETSII)

USC Information Sciences Institute

Sponsored by
Defense Advanced Research Projects Agency
DARPA Order No. J178, J804, J957

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the U.S. Government.

AIR FORCE RESEARCH LABORATORY
INFORMATION DIRECTORATE
ROME RESEARCH SITE
ROME, NEW YORK

STINFO FINAL REPORT

This report has been reviewed by the Air Force Research Laboratory, Information Directorate, Public Affairs Office (IFOIPA) and is releasable to the National Technical Information Service (NTIS). At NTIS it will be releasable to the general public, including foreign nations.

AFRL-IF-RS-TR-2004-224 has been reviewed and is approved for publication

APPROVED: /s/
BRADLEY J. HARNISH
Project Engineer

FOR THE DIRECTOR: /s/
WARREN H. DEBANY, JR.
Technical Advisor
Information Grid Division
Information Directorate

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 074-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE July 2004	3. REPORT TYPE AND DATES COVERED FINAL Apr 00 – Dec 03	
4. TITLE AND SUBTITLE RESEARCH NETWORKS AND TECHNOLOGY MIGRATION (RESNETSII)			5. FUNDING NUMBERS C - F30602-00-C-0151 PE - 62301E PR - J178 TA - 03 WU - 01	
6. AUTHOR(S) Tom Lehman				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) USC Information Sciences Institute 4676 Admiralty Way Marina Del Rey CA 90292-6695			8. PERFORMING ORGANIZATION REPORT NUMBER N/A	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Defense Advanced Research Projects Agency 3701 North Fairfax Drive Arlington VA 22203-1714			10. SPONSORING / MONITORING AGENCY REPORT NUMBER AFRLIF-RS-TR-2004-224	
11. SUPPLEMENTARY NOTES DARPA Program Manager: Gary Koob AFRL Project Engineer: Bradley J. Harnish/IFGA/(315) 330-1884 Bradley.Harnish@rl.af.mil				
12a. DISTRIBUTION / AVAILABILITY STATEMENT APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.			12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 Words) This final report covers activities completed as part of the Research Networks and Technology Migration (RESNETSII) project. The RESNETSII project was part of a multi-agency collaboration conducting research into advanced network architectures and high performance networked applications. One of DARPA's key roles in this effort was the sponsorship of multiple regional high performance networks as well as associated network and application research. These regional networks included the Advanced Technology and Demonstration Network (ATDNET), National Transparent Optical Network (NTON), Boston Optical South Network (BoSSNET), and the Optical Network for Regional Access with Multi-wavelength Protocols (ONRAMP). The RESNETSII project was responsible for several key research and integration functions associated with the NGI objectives and regional network activities. This included: 1) design and implementation of a wide area infrastructure which integrated the regional networks, 2) development and operation of a network operations center for the integrated network, 3) coordination and support for network and application research activities, and 4) conduct of independent research activities. Details of these activities are presented in this technical report. Observations, conclusions, and opinions are offered with regard to the value and utility of such testbed programs. A key note for potential future programs is that networks and network technology diversity has greatly increased since the initial formation of the described research infrastructures. A recommendation is made that further investment in the areas of intelligent Network Management, Resource Provisioning, and Network Control will be required to maximize benefits of future testbed efforts.				
14. SUBJECT TERMS RESNETS, Network, Protocols, IETF, CAIRN, NGI			15. NUMBER OF PAGES 36	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT UL	

Table of Contents

1.	Introduction	1
2.	Task Objectives	1
2.1.	Regional Network Integration	1
2.2.	Wide Area Network Design and Implementation	2
2.3.	Network Operations Center Development and Operation.....	3
2.4.	Peering/Integration with other agency Research networks	3
2.5.	Research Community and Activities	3
2.6.	Advanced Network and Application Research.....	3
3.	Technical Problems	3
4.	General Methodology	4
5.	Technical Results.....	5
5.1.	Regional Network Integration	5
5.2.	Wide Area Network Design and Implementation	7
5.3.	Peering/Integration with other agency Research Networks.....	8
5.4.	Network Operations Center Development and Operation.....	8
5.5.	Research Community and Activities	13
5.6.	Advanced Network and Application Research.....	25
6.	Important Findings and Conclusions.....	30
7.	Implications for Further Research	31
8.	References	32

List of Figures

FIGURE 1	DARPA SUPERNET RESEARCH NETWORK	6
FIGURE 2	NETWORK STATUS MONITORING TOOL	10
FIGURES 3A	REAL TIME TRAFFIC METER	10
FIGURES 3B	REAL TIME TRAFFIC METER.....	11
FIGURE 4	MRTG YEARLY TRAFFIC SUMMARY (DAILY AVERAGE) FOR SUPERNET BORDER ROUTER	11
FIGURE 5	WEB CALENDAR SYSTEM	12
FIGURE 6A	MATISSE PROJECT ARCHITECTURE USING DARPA SUPERNET.....	16
FIGURE 6B	MATISSE PROJECT ARCHITECTURE USING DARPA SUPERNET	16
FIGURE 7	ACAIRN INFRASTRUCTURE AS AN OVERLAY ON DARPA SUPERNET.....	18
FIGURE 8	DARPA NGI PRINCIPAL INVESTIGATORS MEETING.....	22
FIGURE 9	OPTICAL FIBER CONFERENCE 2000 NETWORK CONFIGURATION.....	23
FIGURE 10	SUPERCOMPUTING 2001 NATIONAL COORDINATION OFFICE (NCO) BOOTH.....	24
FIGURE 11	PACKET LOSS AND REORDERING RESEARCH INFRASTRUCTURE	27
FIGURE 12	EYE DIAGRAMS OF GIGABIT ETHERNET CHANNEL AT MULTIPLE TRANSMISSION DISTANCES.....	29

List of Tables

TABLE 1	SUPERNET RESEARCH ACTIVITIES	13
TABLE 2	ACAIRN SITES AND PEERING NETWORKS.....	18
TABLE 3	ACAIRN RESEARCH ACTIVITIES	19

1. Introduction

This final report covers activities completed as part of the Research Networks and Technology Migration (RESNETSII) project. The RESNETSII project was part of a multi-agency collaboration conducting research into advanced network architectures and high performance networked applications. This multi-agency collaboration was known as the Next Generation Internet (NGI) project. The NGI project objectives included the following: i)advance research, development, and experimentation in next generation network technologies, ii)develop an advanced wide area network testbed, emphasizing end-to-end performance, iii)develop and demonstrate advanced applications.

One of DARPA's key roles in this effort was the sponsorship of multiple regional high performance networks as well as associated network and application research. These regional networks included the Advanced Technology and Demonstration Network (ATDnet), National Transparent Optical Network (NTON), Boston Optical South Network (BoSSNET), and the Optical Network for Regional Access with Multi-wavelength Protocols (ONRAMP). The RESNETSII project was responsible for several key research and integration functions associated with the NGI objectives and regional network activities. This included i)design and implementation of a wide area infrastructure which integrated the regional networks, ii)development and operation of a network operations center for the integrated network, iii)coordination and support for network and application research activities, iv)conduct of independent research activities.

The RESNETS II project and other NGI programs continued the contribution made by DARPA in creating the initial Internet architecture. These programs provided for basic research into broad range of networking topics including physical layer optical technologies, network security, network architecture, and application development. This research contributed in many ways to advancing the state of the art in networking and included the migration of NGI technologies to the government and commercial deployment.

This report describes these activities and is structured as follows: Section 2 describes the task objectives, Section 3 describes the technical problems, Section 4 describes the general methodology, Section 5 describes the technical results, Section 6 describes the important findings and conclusions, Section 7 describes the implications for future research activities, and Section 8 contains a reference list.

2. Task Objectives

The objective of the RESNETSII project was to develop, maintain, and support a network testbed to allow the network research community to further their research goals in a broad range of network protocol areas.

The RESNETSII project had multiple objectives relating to advanced network and high performance application research in accordance with the overall objectives of the NGI program. Many of the activities were interrelated and highly dependent on each other. Specific task objectives were:

- Regional Network Integration
- Wide Area Network Design and Implementation
- Peering/Integration with other agency research networks
- Network Operations Center Development and Operation
- Research Community and Activities
- Advanced Network and Application Research

2.1. Regional Network Integration

A key task objective for the RESNETSII project was to develop infrastructure, software, and technologies to integrate the regional network infrastructures funded by DARPA. The goal was to provide an environment such that researchers could construct and conduct experiments which spanned multiple of these regional networks. These regional research network infrastructures funded by DARPA are briefly described below:

Advanced Technology and Demonstration Network (ATDnet)

ATDnet is a high performance networking testbed in the Washington D.C. area. Established by DARPA to enable collaboration among Defense and other Federal agencies, a primary goal of ATDnet was to serve as an experimental platform for diverse network research and demonstration initiatives. ATDnet activities included early deployment of Asynchronous Transfer Mode (ATM), Synchronous Optical Network (SONET), and Wave Division Multiplexing (WDM) technologies. During the duration of the RESNETS project, ATDnet deployed an optical network based on prototype optical equipment from Lucent and Tellium. This network deployment was accomplished as part of the Multiwavelength Optical Networking (MONET) consortium whose members included AT&T, Bell Atlantic, BellSouth, Lucent Technologies, Pacific Telesis, SBC/TRI, and Telcordia.

National Transparent Optical Network (NTON)

NTON was a high performance network infrastructure in California spanning from San Francisco to Los Angeles. NTON was developed by the NTON Consortium (NTONC) whose membership consisted of Nortel Networks, GST Telecom, Sprint Communications, and Lawrence Livermore National Laboratory.

NTON was a Wavelength Division Multiplexed network deployed using commercial fiber. NTON linked multiple government, research and private sector labs and provided multiple network services including OC192/OC48 SONET, ATM, IP, and WDM connections. The NTON network utilized Nortel equipment for the SONET and WDM service, FORE Systems for ATM capabilities, and AVICI routers for IP service.

NTON was a research infrastructure for developing and showcasing high-bandwidth applications and for field-testing emerging technologies -- including new optical devices, new protocols, and new management paradigms.

BoSton-South NETwork (BoSSNET)

BoSSNET is wide area dark fiber based network developed and managed by the Massachusetts Institute of Technology/Lincoln Laboratory (MIT/LL). The testbed utilizes an embedded dark fiber plant in the eastern United States that is to be supplied by Qwest Communications which runs between Boston, Massachusetts and Washington D.C. The goals of this system, which spans over 1000 route miles, is to extend the state-of-the-art in non-repeated (optically amplified, without electronic conversions) optical network systems. This infrastructure was used as part of the NGI program for both network layer research as well as advanced application research thru collaborations with the MIT Haystack Observatory and MIT main campus.

Optical Network for Regional Access with Multi-wavelength Protocols (ONRAMP).

The Optical Network for Regional Access using Multi-wavelength Protocols (ONRAMP) was a consortium sponsored by DARPA. Its mission was to develop architectures, protocols, and algorithms for wavelength division multiplexing (WDM)-based regional access networks that will effectively support the NGI. A reconfigurable WDM test bed was being built at MIT in Boston, Massachusetts to demonstrate some of the key thrusts of the consortium, including dynamic service provisioning and optical flow switching, service protection in the optical domain, medium access control protocols, and network control and management geared for the efficient transport of Internet traffic over WDM networks. The ONRAMP test bed consisted of a feeder network connected via access nodes to distribution networks on which the end users reside. ONRAMP network reconfiguration is enabled by access nodes that contain both optical and electronic switching components, allowing data traffic to be routed all-optically through the network or to be switched and aggregated by electronic Internet Protocol (IP) routers.

An important part of this integration was the wide area network infrastructure which is defined in the following section.

2.2. Wide Area Network Design and Implementation

A key activity of the RESNETSII project was integrating the regional networks via high speed wide area network access. The High Speed Connection Consortium (HSCC) was separately funded by DARPA to

contract with Qwest Communications for OC48 connections at telecommunications facilities in Los Angeles, California, Pittsburgh, Pennsylvania, and Washington D.C. RESNETSII project activities included designing and implementing network facilities to connect regional networks to the national infrastructure as well as interconnecting regional infrastructures across the wide area infrastructure. The combination of the regional networks integrated into the HSCC wide area infrastructure was often referred to as the DARPA SuperNet[1] network.

2.3. Network Operations Center Development and Operation

A Network Operations Center (NOC) was organized and operated in order to maintain and monitor the network as well as facilitate research activities. Primary NOC duties included providing assistance to researchers, troubleshooting, maintaining and upgrading links and routers, maintaining mailing lists and web pages, maintaining code distributions, and expanding the network to new sites.

The NOC also provided support at the Principal Investigator meetings and demonstrations. This included networking support for onsite demonstrations as well as high speed connections to the research network to allow individual researchers the ability to access their research facilities from remote meeting sites.

2.4. Peering/Integration with other agency Research networks

The combined network infrastructure, DARPA SuperNet was organized as a single Autonomous System (AS) for purpose of peering with the University Consortium for Advanced Internet Development (UCAID) Internet2 Abilene Network, the Mid-Atlantic Crossroads (MAX), and the Qwest Commercial Internet. This provided access to other national University and Federal Research and Education (R&E) networks including:

- Department of Energy (DOE) Energy Sciences Network (ESNET)
- NASA Research and Education Network (NREN)
- Defense Research and Education Network (DREN)

2.5. Research Community and Activities

The DARPA SuperNet network was utilized to support a wide range of network researchers and research activities. The majority of the research activities were separate programs directed by various Principal Investigators (PIs) who were funded from a wide range of government, education, and commercial organizations. The SuperNet use policy remained very open to researchers interested in conducting experiments and trials on the research network infrastructure. A key goal of the network was to support the networking research community which was actively involved in developing standards, building software, and conducting research in the area of network layer protocols. This included many individuals actively involved with the Internet Engineering Task Force (IETF) and other standard bodies. This allowed research results to be factored into standard bodies' activities. Additional details regarding these activities are presented in the Technical Results section.

2.6. Advanced Network and Application Research

RESNETSII project activities also included direct conduct of and participation in advanced network and high performance application research. This included research into advanced network architectures, end-to-end performance across large bandwidth-delay product networks and distribution of uncompressed high definition television format data across IP networks, and gigabit rate IP security testing.

Additional details on these Task Objectives and the results are presented in the Technical Results section.

3. Technical Problems

Advanced network and application researchers are often limited to analysis, simulations, and lab testing of their research theories and developments. Sometimes, they may have access to a regional network on which to conduct their experiments. However, these regional networks typically are not high performance and do not directly connect to the important research facilities. These research facilities may include large and expensive resources like computational clusters, databases, visualization engines, or science

instruments like radar antennas, radio telescopes, or electronic microscopes. The DARPA funded NGI program sought to address these issues via the development of several regional high performance infrastructures which did provide advanced network services directly to key researchers, research resources, and facilities. An important technical problem the RESNETSII project had to address was how to integrate these regional infrastructures in a manner that would allow and encourage the following:

- access to research facilities from other regions
- enable new experiments and research not possible on a regional level
- encourage and enable new collaborations amongst groups of researchers not previously able to conduct joint experiments

The technical problems and challenges for this project included:

- integration of regional network which had different network technologies and primary research objectives
- isolation from production services
- coordination and scheduling of competing research activities amongst researchers
- monitoring/maintenance of the infrastructure and research activities.
- development of tools and procedures to facilitate researchers use of the infrastructure

The approach the RESNETSII project utilized evolved over the life of the program as technology and researcher objectives progressed. The general methodology is described in the following section.

4. General Methodology

A key objective of the RESNETS II project was to continue the contribution made by DARPA in creating the initial Internet architecture and continuing to provide an environment to allow associated research activities to continue. These programs provided for basic research into broad range of networking topics including physical layer optical technologies, network security, network architecture, and application development. This research contributed in many ways to advancing the state of the art in networking and included the migration of NGI technologies to government and commercial deployment.

The regional networks which constituted the DARPA SuperNet each incorporated different network technologies, had a different primary research focus, and engaged a different set of researchers. This research and network topology diversity was considered a strength, although it did provide some challenges in designing the wide area infrastructure and support tools to allow integration across these regional facilities.

Interoperability across the various NGI research networks was increasingly problematic as the number of networks and networking technologies continue to increase. Our approach was to experiment with different methodologies for integrating network technologies such as Packet Over SONET (POS), ATM, WDM, Gigabit Ethernet, Fast Ethernet, T1, ISDN, ADSL, and wireless technologies. We also wanted to utilize production services when available to leverage the investment in that infrastructure. Our approach toward this was to utilize overlay networks in order to provide researchers with the "dedicated" network environment they sought. An important part of this effort was to develop techniques to adequately isolate the overlays from the underlying infrastructure to ensure no perturbation to production services.

The use of Gigabit Ethernet technologies within Wide Area Optical Networks was identified as an excellent candidate for use in the wide area as well as an inter-network peering technology. We conducted early tests on use of Gigabit Ethernet in the wide area and incorporated that into our designs.

A key part of our methodology was development of SuperNet Point of Presence (SPOP) facilities to the interconnection of the regional networks, as well as provide peering points with other networks. We utilized a mix of technologies including IP tunnels, Multi Protocol Label Switching (MPLS), individual wavelengths in WDM networks, and ATM Virtual Circuits (VC) to combine a physically distributed set of

network elements into what appeared to be a central peering point. Major SPOPs were established in the Washington D.C. and Los Angeles areas.

The general methodology applied by the RESNETS II project was to accept and encourage this regional network diversity and heterogeneity. A flexible set of technologies and tools were utilized to allow the regional networks to be integrated into a unified testbed environment. These included application of routing policy to control traffic flow across the wide area network, development of open source router platforms to allow researchers to construct experiment specific overlay infrastructures, and careful identification and control of external peering points to ensure consistent connectivity to other research infrastructures.

The specific infrastructure architecture, open system host/router platform, and research activities evolved over the course of the program. These specific activities are described in the following section.

5. Technical Results

The objective of the RESNETSII project was to develop, maintain, and support a network testbed to allow the network research community to further their research goals in a broad range of network protocol areas.

The RESNETSII project had multiple objectives relating to advanced network and high performance application research in accordance with the overall objectives of the NGI program. Many of the activities were interrelated and highly dependent on each other. Specific task objectives were:

- Regional Network Integration
- Wide Area Network Design and Implementation
- Peering/Integration with other agency research networks
- Network Operations Center Development and Operation
- Research Community and Activities
- Advanced Network and Application Research

Technical results resulting from the completion of these tasks are presented below.

5.1. Regional Network Integration

A key task objective for the RESNETSII project was to develop infrastructure, software, and technologies to integrate the regional network infrastructures funded by DARPA. The goal was to provide an environment such that researchers could construct and conduct experiments which spanned multiple regional networks. The regional networks included the following:

- Advanced Technology and Demonstration Network (ATDnet)
- National Transparent Optical Network (NTON)
- Boston Optical South Network (BoSSNET)
- Optical Network for Regional Access with Multi-wavelength Protocols (ONRAMP)

A high speed wide area network was constructed to interconnect these regional infrastructures. The High Speed Connection Consortium (HSCC) was separately funded by DARPA to contract with Qwest Communications for OC48 connections at telecommunications facilities in Los Angeles, California, Pittsburgh, Pennsylvania, and Washington D.C. RESNETSII project activities included designing and implementing network facilities to connect regional networks to the national infrastructure as well as interconnecting regional infrastructures across the wide area infrastructure. The combination of the regional networks integrated into the HSCC wide area infrastructure was often referred to as the DARPA SuperNet network. A map of the SuperNet infrastructure is shown in Figure 1.

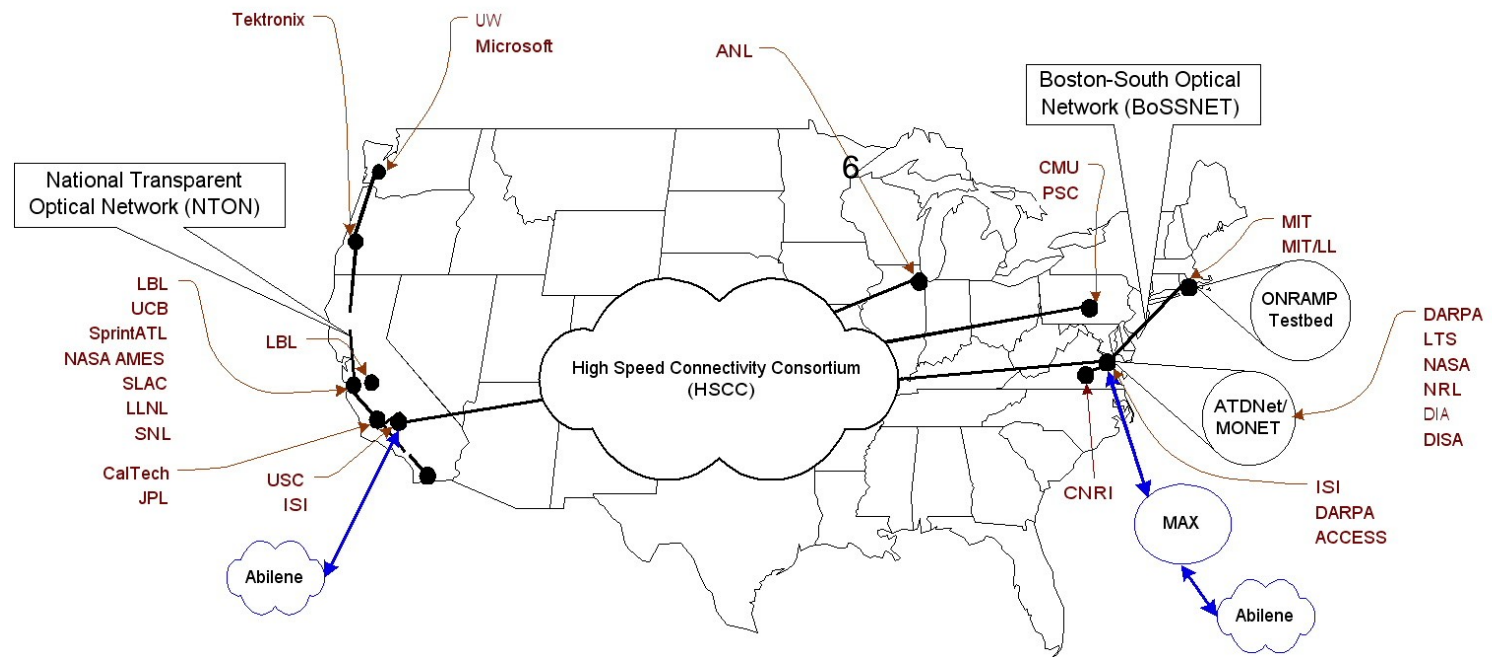


Figure 1. DARPA SuperNet Research Network

5.2. Wide Area Network Design and Implementation

A key activity of the RESNETSII project was integrating the regional networks via high speed wide area network access. The High Speed Connection Consortium (HSCC) was separately funded by DARPA to contract with Qwest Communications for OC48 connections at telecommunications facilities in Los Angeles, California, Pittsburgh, Pennsylvania, and Washington D.C. RESNETSII project activities included designing and implementing network facilities to connect regional networks to the national infrastructure as well as interconnecting regional infrastructures across the wide area infrastructure. The combination of the regional networks integrated into the HSCC wide area infrastructure was often referred to as the DARPA SuperNet network.

The HSCC provided access to the Qwest commercial internet backbone. This access was provided in the form of OC48 SONET interfaces into IP routers at Washington D.C., Los Angeles, California, and Pittsburgh Pennsylvania. Routing across the Qwest commercial backbone was controlled via use of specific Border Gateway Protocol (BGP) communities and route advertisements. The result was that the interconnection between the regional networks enjoyed relatively high speed connection. Thru the use of IP in IP tunnels across the Qwest backbone, we were able to control routing between the regional networks in manners to support individual experiment and research objectives.

An important issue was that traffic flowing across the HSCC wide area infrastructure was routed along with the commercial internet traffic. That is, while the BGP route advertisements included distinguishing community strings to limit reachability to the DARPA research facility involved, the forwarding path for the data packets was mixed in with the general internet traffic. Since the research network connections were directly connected to Qwest backbone routers, performance across the HSCC network was excellent. The Qwest backbone was a diverse, over provisioned OC48 network which evolved to a 10 Gigabit/s backbone by the end of this DARPA research network program.

While the wide area connections did not provide for dedicated physical paths between the research networks, the interaction with actual internet traffic did provide some interesting opportunities for performance testing under real world network conditions, which are described in more detail in the Advanced Network and Application Research section below.

When required, IP-in-IP tunnels and overlay architectures were utilized to provide the researchers with a network routing and protocol environment which was dedicated. That is, while there were not hard guarantees on amount of bandwidth available, the overlay design allowed the research network to run its own instance of routing protocols which provided the researchers with a "network they could break".

Washington D.C. SuperNet Point of Presence (SPOP)

A Washington D.C SPOP was implemented and provided an interconnect for BoSSNET, ATDNet, HSCC, Mid-Atlantic Crossroads (MAX), Abilene, and ISI East. This POP also facilitated the traffic engineering options for data flows across the HSCC backbone. This included the ability to create MPLS and Multicast tunnels. The Abilene peering allowed collaboration between SuperNet and Internet2 researchers. The SPOP was a distributed design and not all peering network elements were located in same facility. Use of MPLS, and Gigabit Ethernet over WDM were utilized in a (at the time) novel fashion to provide for a "virtual" peering facility.

The SuperNet to Abilene peering was utilized as an early implementation and research infrastructure for future peering architectures. We investigated advanced networking peering architectures such as optical cross-connects, multi-hop BGP peering, layer two router switching, and MPLS LSPs. We implemented an early version of a distributed peering architecture in the form of a layer 2 tunnel across MAX to provide for a direct BGP peering between SuperNet and Abilene. The layer2 tunnel was built from MPLS LSPs and Juniper router Circuit CrossConnect (CCC) technology and demonstrated how to implement a distributed peering topology.

Los Angeles SuperNet Point of Presence (SPOP)

A Los Angeles Area SPOP was implemented and provided an interconnect for NTON, HSCC, and Abilene at OC48 rates and USC/ISI via gigabit ethernet. This POP also facilitates the traffic engineering options for data flows across the HSCC backbone. This includes the ability to create MPLS and Multicast tunnels. The Abilene peering allowed collaboration between SuperNet and Internet2 researchers.

5.3. Peering/Integration with other agency Research Networks

Peering with other national research and education networks was utilized to broaden the reach of both DARPA sponsored researchers as well as other organization sponsored research initiatives. The results of these collaborations were demonstrated at various conferences and Principal Investigator meetings. As an example, the 2001 DARPA NGI Principal Investigators meeting included the following projects which utilized the SuperNet to Abilene peering to demonstrate their research results:

- Distributed Classroom , Microsoft
- TeleImmersion , UNC
- Pegasus, Drexel

Additional details on this meeting including a diagram of the network and peering infrastructure are provided in Section 5.5.7.1 of this report.

Peering between DARPA research networks was also used to support Abilene activities including the March 2001 Internet2 Member meeting where the following advanced applications were demonstrated:

- Stereoscopic Rendered Images and Video Streaming with Real-time Compression Methods
- Telepresence in the Operating Room Utilizing IP Video

5.4. Network Operations Center Development and Operation

A Network Operations Center (NOC) was organized and operated in order to maintain and monitor the network as well as facilitate research activities. Primary NOC duties included providing assistance to researchers, troubleshooting, maintaining and upgrading links and routers, maintaining mailing lists and web pages, maintaining code distributions, and expanding the network to new sites.

The responsibilities of the NOC included the following:

- Network design, build, maintenance, and monitoring
 - Network link support
 - Network routing and peering
 - Network backbone monitoring and maintenance
- Community coordination and researcher assistance
 - Maintain Web Site
 - Maintain mailing lists
 - Maintain SuperNet Calendar system for experiment scheduling
 - Video Teleconferencing Support
- Open Source Performance Host and Router build, test, and deployment
 - software specification and configuration
 - hardware specification and configuration
 - system build, test, and deployment
 - Software repository for network specific configurations

Additional details on these efforts are described below.

5.4.1. Network design, build, maintenance, and monitoring

The research network evolved throughout the years as the regional networks were introduced, changed configurations, or were turned down. The NOC developed the architectures and designs for these networks and the associated transitions. In addition, the NOC developed the layer 3 routing architecture,

protocol configurations, and peering agreements with other networks. The IP protocols include IPv4, IPv6, and multicast.

The NOC responsibilities also included monitoring of the network backbone and router nodes. The tools utilized provided status on link and node availability, bandwidth utilization, operating system configurations, cpu utilization, available disk resources, and other data.

Various tools were utilized during the duration of the project and included the following:

Netarchive Data Plotter[2]

The Information & Telecommunication Technology Center at the University of Kansas developed the Netarchive Data Plotter architecture. The network archive architecture includes a configuration data base, time series database, traffic and connectivity information collectors, and various plot and information summary utilities. Included among the types of information collected are throughput measurements based on switch cell and router packet counts, as well as connectivity and round-trip time information based on ping.

SuperNet Developed ToolSet

Network monitoring tools developed as part of the RESNETS II Program were utilized to monitor the SuperNet testbed. These tools allowed for customization of monitoring data and web based presentation formats. For the SuperNet network, these tools were configured to allow researchers to view network status, real-time usage, and historical usage. Screen shots of these tools as configured for SuperNet are shown in Figures 2 and 3.

MRTG[3]

The Multi Router Traffic Grapher (MRTG) is a tool to monitor the traffic load on network-links. MRTG generates HTML pages containing GIF images which provide a live visual representation of this traffic. MRTG is based on Perl and C and works under UNIX and Windows NT. Screen shots of the MRTG historical data as configured for SuperNet is shown at Figure 4.

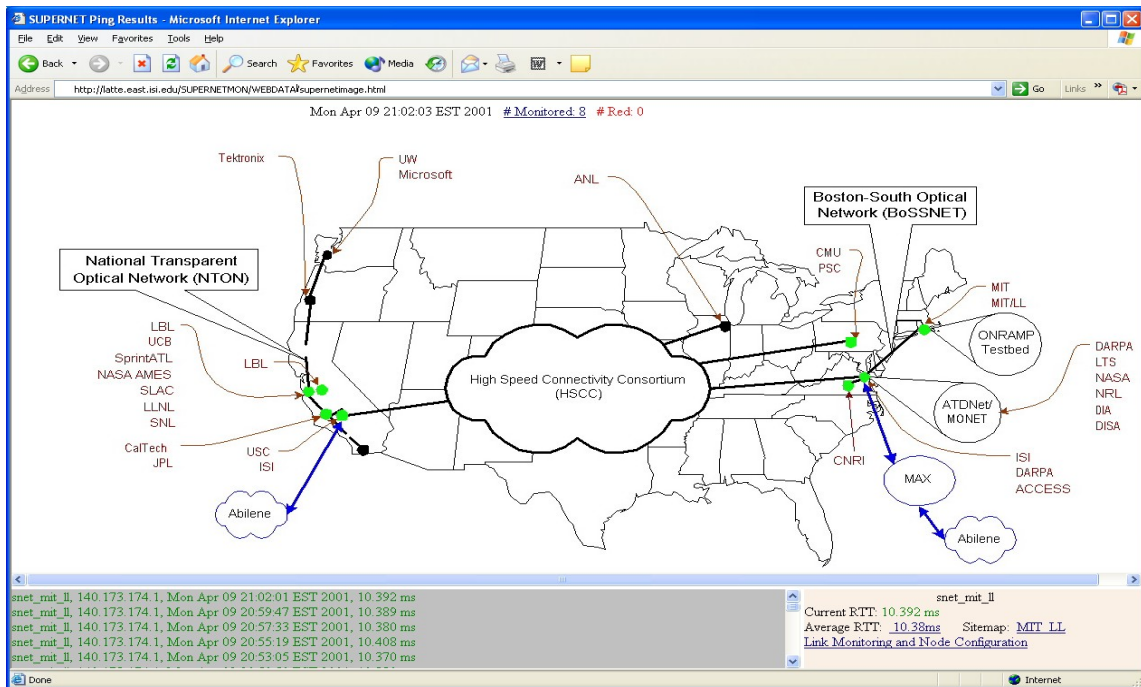
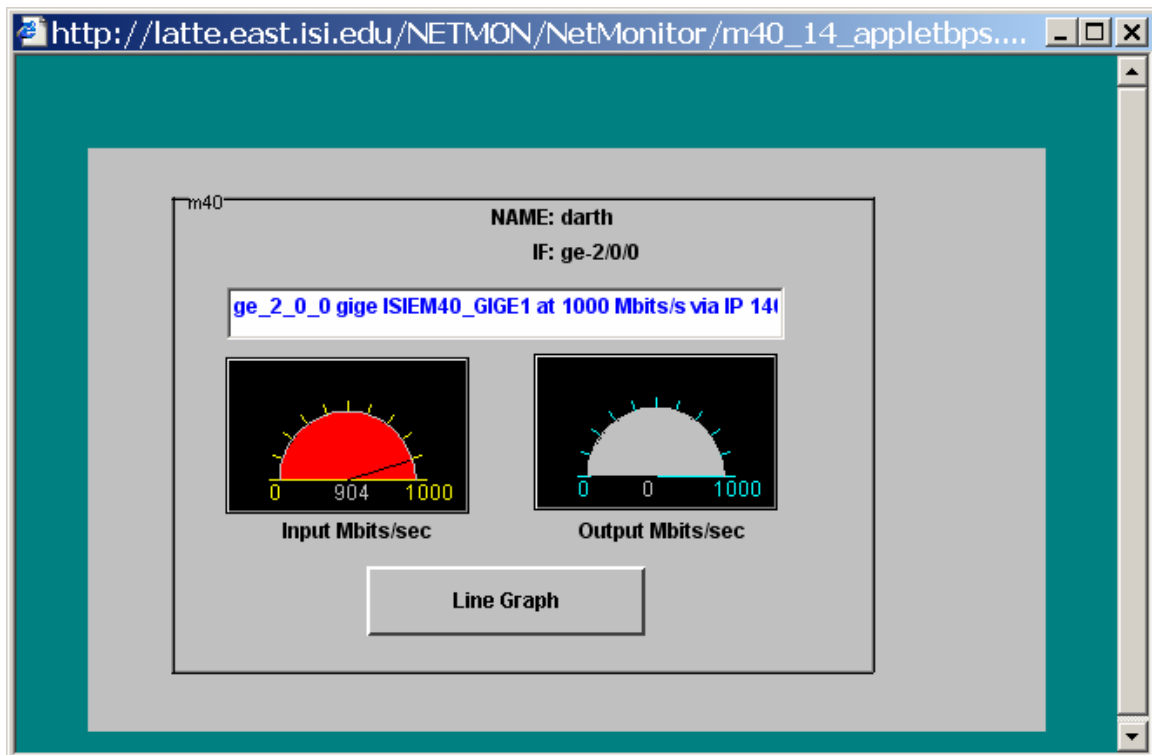
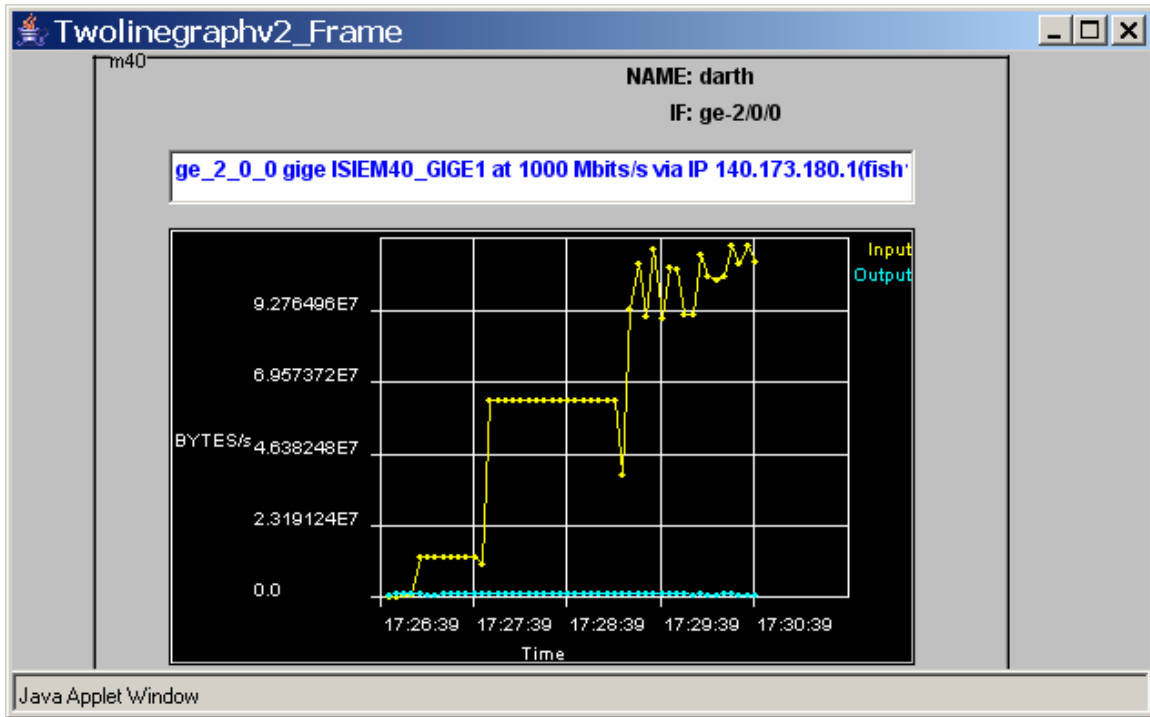


Figure 2 Network Status Monitoring Tool



Figures 3a Real Time Traffic Meter



Figures 3b Real Time Traffic Meter

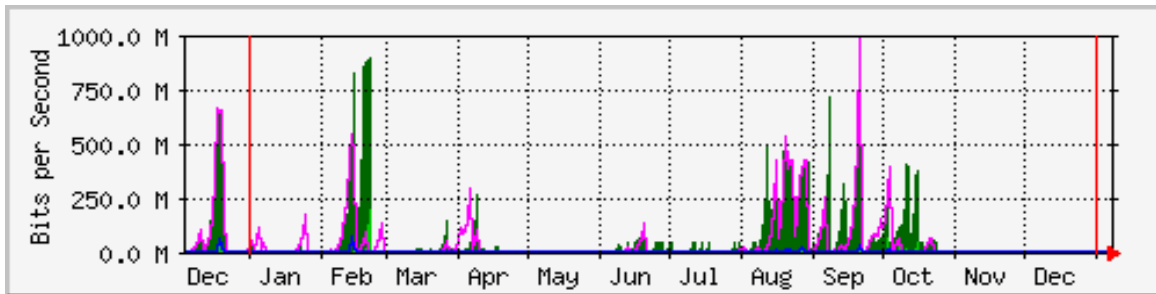


Figure 4 MRTG Yearly Traffic Summary (Daily Average) for SuperNet Border Router

5.4.2. Community coordination and researcher assistance

The NOC provided assistance and tools for the research community to schedule experiments, reserve network resources, and collaborate via video teleconference systems. The NOC also provided assistance in reconfiguration of the network as required for specific experiment and research objectives. This included maintenance of a testbed web site. A Web Calendar system was utilized to allow researchers to schedule parts or the entire network for experimental activities. A screen shot of the calendar system interface is shown in .Figure 5.

5.4.3. Open Source Performance Host and Router

Configuration of the network for specific researcher requests often required the building of open source router and performance hosts to satisfy specific experiment requirements. The NOC developed

configurations and designs for these open source router platforms and performance host. The performance hosts were specially configured hosts tuned to perform well when operated across high bandwidth delay product networks. The open source router platforms allowed specific configurations of routing protocols and architectures to be implemented as an overlay architecture on top of the research network. These open source platforms were built using FreeBSD and Linux operating systems.

The NOC was responsible for the building, testing, and deployment of these performance host and router platforms. This included configuring kernels specific to the research objectives, porting and configuration of driver and application code, and generation of scripts to update router systems.

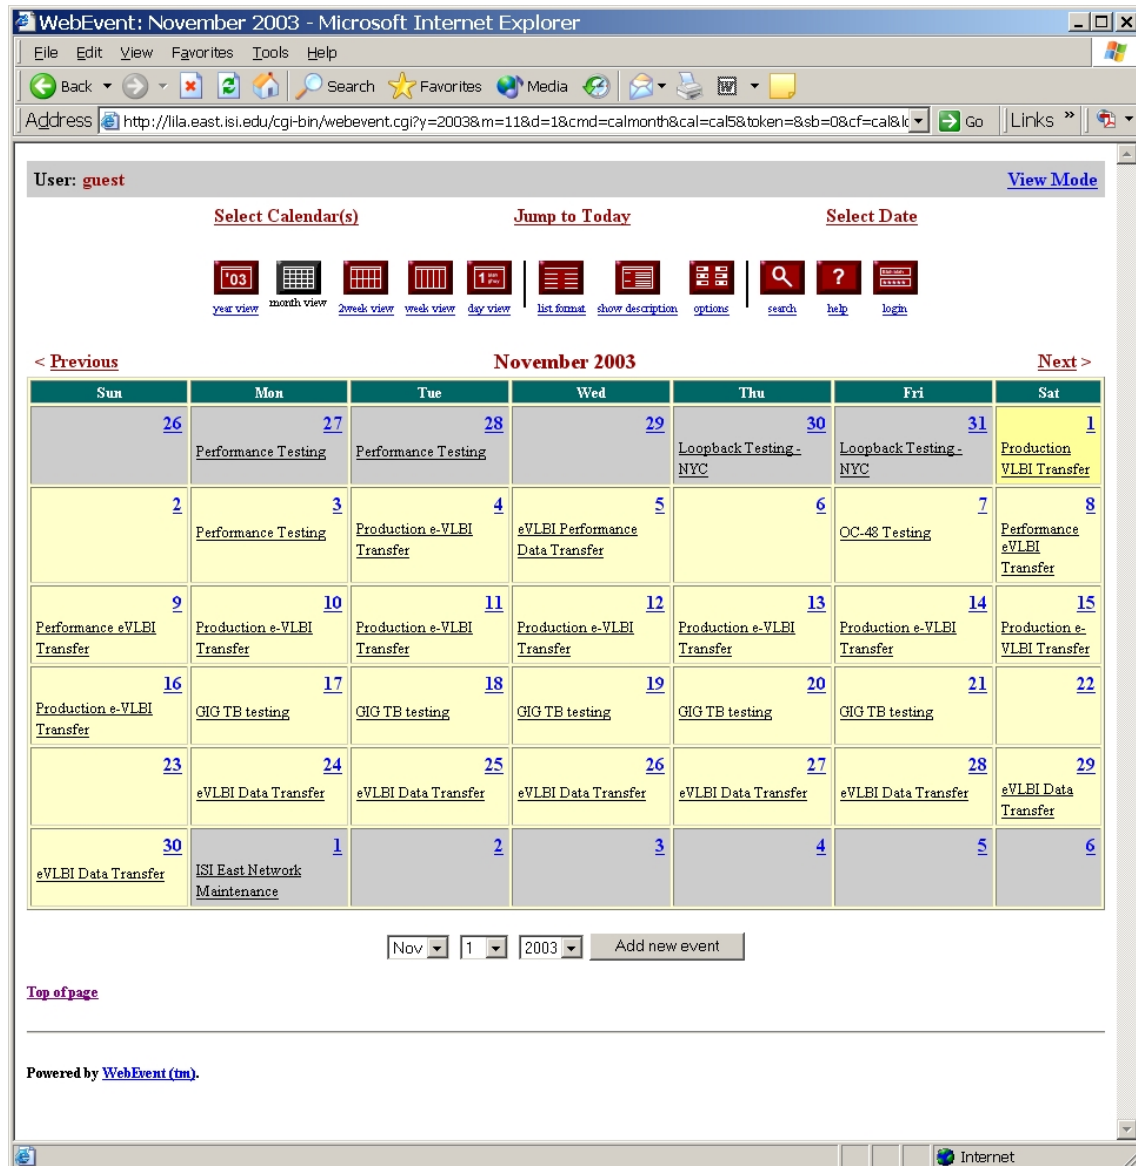


Figure 5 Web Calendar System

5.5. Research Community and Activities

The DARPA SuperNet network maintained a very open use policy in terms of allowing individual researchers to decide which types of experiments to conduct. The approach taken was to focus on connecting a facility and a group of researchers and allowing the collaborations to determine the research directions. Research projects which utilized the DARPA SuperNet network for research and experimentation activities are listed in Table 1. This list is likely missing some of the organizations and research activities which were conducted on the testbed. This is no reflection on those activities, but since the actual research principals were funded via many sources, the results did not always make it to the operations center for historical documentation. The majority of the research activities involved the cooperation and collaboration of many people and organizations. However, for brevity purposes, only those primarily responsible for the research agenda and execution are listed in the Primary Organizations column.

Table 1 SuperNet Research Activities

Research Project	Research Topics	Primary Organizations
Very Long Baseline Interferometry (VLBI)	Application of high speed networks to real time correlation of radio telescope data for GPS calibration and other scientific applications	MIT Haystack Observatory, USC/ISI, NASA, USNO
WideBand Networked Sensors Program (WNS)	inter-domain multicast routing	MIT/LL
Remote Media Immersion (RMI)	capturing, streaming, and rendering high-resolution, big-screen digital video and multichannel audio	USC, ISI
MATISSE	Distributed application for collaborative MEMS system design, implementation, and test	CMU, CNRI, ISI, LBNL, MIT, Sarnoff, UCB
Distributed-Parallel Storage System (DPSS)	high-performance, distributed-parallel data storage	LBNL
Collaborative, Operational Virtual Exploitation Team (COVET)	Collaborative application development of military applications	Veridian, Internet2, NASA GSFC, US Army Corps of Engineers
Gigabit to the Desktop	deployment and gigabit capable end systems to desktop	USC/ISI, UC Berkeley, University of Washington, CMU, MIT
High Performance Local Area Networks	research into optical layer impairments compensation to realize 40 Gb/s channel rates	UMBC
NGI Multicast Applications and Architecture (NMAA)	large scale high definition internet multimedia and conferencing systems	USC/ISI
Uncompressed HDTV	uncompressed HTDV collaborative internet applications	USC/ISI
Advanced Collaborative Advanced Interagency Research Network (ACAIRN)	overlay network dedicated to network layer research	multiple, see expanded description below

Research Project	Research Topics	Primary Organizations
Supercomputing Conference	Active participation in multiple conferences, experiments, and demos	multiple, see expanded description below
Digital Earth	technologies to enabling a massive, scalable, and open model of the planet which is networked for access which scales up to millions of user	SRI
Internet2 Land Speed Record	Set Land Speed Record in 2000	USC/ISI, University of Washington, Microsoft, Qwest
Extended Gigabit LAN Services	research into using gigabit ethernet as a wide area network protocol across optical networks	USC/ISI, Telcordia

While there were many research and experimentation activities during the duration of the testbed, several are described in more detail below.

5.5.1. Very Long Baseline Interferometry (VLBI)

Very-Long-Baseline Interferometry (VLBI) has been used by radio astronomers for more than 30 years as one of the most powerful techniques for studying objects in the universe at ultra-high resolution and measuring earth motions with ultra-high accuracy [4]. VLBI allows images of distant radio sources to be made with resolutions of tens of microarcseconds, far better than any optical telescope. VLBI also provides a stable inertial reference frame formed by distant quasars to study the motions of the Earth in space with exquisite precision, revealing much information about both surface and internal motions of the Earth system, including interactions with the dynamic motions of the atmosphere and oceans.

VLBI combines data simultaneously acquired from a global array of up to ~20 radio telescopes to create a single coherent instrument. Traditionally, VLBI data are collected at data rates close to ~1 Gbps on magnetic tapes or disks that are shipped to a central site for correlation processing. This laborious and expensive data-collection and transport process now has the possibility of being replaced by modern global high-speed networks, potentially enabling important new capabilities, real-time data correlation and analysis and scientific returns. By nature, VLBI data are digital representations of the analog signal arriving at a radio telescope. Almost always this is white Gaussian noise sampled at the Nyquist rate so that each sample is independent and the data are uncompressible.

The transmission of VLBI data via high-speed network is dubbed 'e-VLBI'. In the fall of 2002, several ~1 Gbps e-VLBI demonstration observations [5] were performed using antennas in Westford, MA and Greenbelt, MD with correlation at Haystack Observatory in Westford, MA. These experiments were supported by DARPA and NASA/GSFC, with the close cooperation of ISI-E and MAX.

This research and experimentation was conducted over DARPA SuperNet and was instrumental in advancing the state of the art for eVLBI. This work has continued even after the end of DARPA support, with the United States Naval Observatory (USNO) regularly using these procedures and infrastructure to collect eVLBI data as part of their ongoing testing efforts to increase their operational correlation data quality and response time.

5.5.2. Wideband Networked Sensors (WNS)

The DARPA-sponsored Wideband Networked Sensors (WNS) Program conducted research and development for a future vision of wideband distributed sensor systems serving sophisticated military needs. The first major goal of the program was a demonstration of radar data fusion, exploiting the multi-Gbps transmission capabilities of the BoSSNET. The goals of the WNS program were to use BoSSNET to enable real time remote processing, fusion, and distribution of radar data and demonstrate its application to

advanced ultra-wideband communication-intensive military applications such as fusion of multi-source data to support C3 I, battle management, and target identification. As part of the RESNETS II project, we supported MIT/LL in these testing and demonstration efforts.

5.5.3. MATISSE

MEMS co-locate, on single devices, sensing, actuating, and control functions along with computing and communications. MEMS build on micro-electronic manufacturing but are much more complicated because their fabrication involves many specialized steps, and testing must be performed on mechanical as well as electronic operation. Specialized platforms generate massive amounts of data that must be processed, stored, and visualized, and often compared with tests on antecedent devices. Typically, a device designer from industry or academia will travel to a site that hosts one of the test platforms, may stay at the site for a week or more to ensure that the tests have been set up correctly and will then return home to wait for the test results, which may arrive in the form of a written report several weeks later.

The purpose of the MATISSE project was to take advantage of the capabilities of a high -performance network, and greatly reduce the steps outlined above. Data, originating at the source such as a MEMS test platform, is sent first to a network-based, high-speed storage cache, the central component of the architecture. From the cache, the data can be sent to an archive for long-term storage, and/or to a computing facility for processing. The processed data are sent to the user (directly, or via the cache) and/or to the archive. Previously obtained raw or processed data can be retrieved from the archive and re-examined, re-processed, or compared with new results. An important point to note is that the five architectural elements can be located anywhere because they are linked by a network. Moreover, if computing resources, for example, are available at several sites, then the system should be able to select the one that will give best performance (or lowest price). Further, there is no limit to the size of the dataset that can be handled, so long as the cache can temporarily store the volume of data currently needed. The Matisse Project implemented this architecture for the MEMS R&D community.

The Matisse project was a joint collaboration between University of California Berkeley, Lawrence Berkeley Laboratory (LBL), Carnegie Mellon University (CMU), Massachusetts Institute of Technology (MIT), and Corporation for National Research Initiatives (CNRI). The Matisse project utilized the DARPA SuperNet to develop and test their system. Some of the results of Matisse experiments over SuperNet were documented in [6]. Figure 6 depicts their architecture.

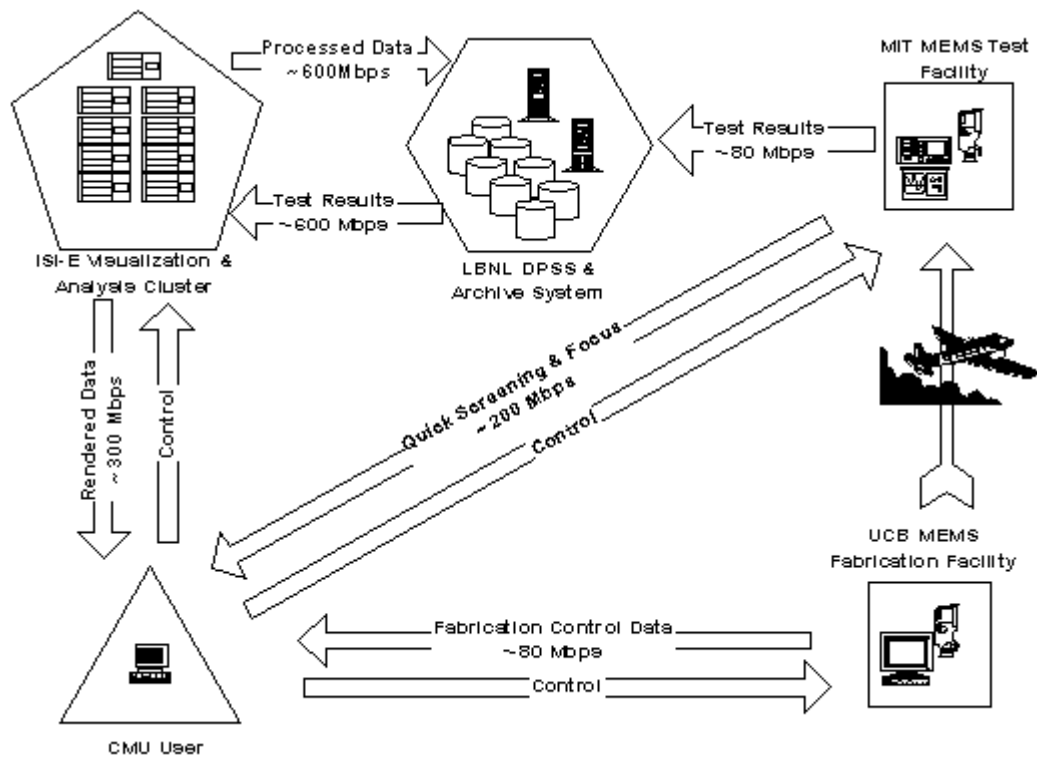


Figure 6a Matisse Project Architecture Using DARPA SuperNet

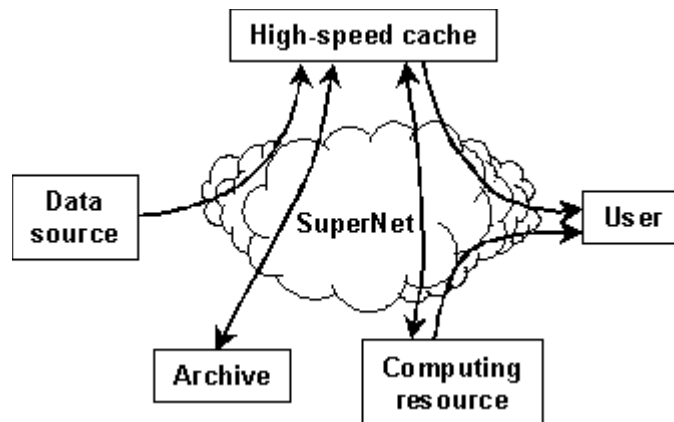


Figure 6b Matisse Project Architecture using DARPA SuperNet

5.5.4. ACAIRN Research Network

The Advanced Collaborative Advanced Interagency Research Network (ACAIRN) was a testbed dedicated to the use of network researchers. Network researchers are interested in experimenting with the network and transport layers on a broad range of topics including routing protocol stability and enhancement, quality of service, multicasting, Internet protocol (IP) version 6, management and infrastructure protocols such as SNMP and NTP, and IP security. There is a large networking research community which is actively involved in developing standards, building software, and conducting research in the area of network layer protocols. Individuals involved in this community are from various government, educational, and commercial institutions and are typically involved with the Internet Engineering Task

Force (IETF) and other standard bodies to advance the state of network technologies. A key objective of the ACAIRN projects was to enable greater collaboration and experimentation amongst the network research community.

The ACAIRN activities were a follow on to the other DARPA research networks including the DARTnet (DARPA Research Testbed network) and Rational ATM Internet Service Environment (RAISE) testbeds. These testbeds were very active in many areas of internet protocol and architecture development.

The technologies and protocols that network researchers are interested in experimenting with are the very technologies that allow the network stability and connectivity to be maintained. This is quite different from application research, where researches require that the network infrastructure remain stable with very little down time. The result is that network research activities require some unique infrastructure components, architecture considerations, and services as compared to that required for applications research. These include such things as open system router protocol stacks, mechanisms to isolate network research activities from “production” network operations, and separate operations management.

In support of the ACAIRN research testbed, the RESNETSII project was responsible for the design and implementation of the wide area network infrastructure over which the ACAIRN testbed was constructed. The requirements included presenting a network infrastructure to the ACAIRN researchers in a form where they could experiment with the network layer technologies and protocols. That is, they needed an experimental infrastructure in the form of a “network they can break”. This was provided via the DARPA SuperNet infrastructure by the use of various overlay technologies. A combination of IP in IP tunnels, and MultiProtocol Label Switch (MPLS) Label Switched Path (LSP) technologies were utilized to develop an infrastructure where the ACAIRN network was run as an overlay on top of SuperNet. This allowed the network researchers to experiment with new and modified network protocols without impacting the operations of the underlying SuperNet, over which other research and experiments were being conducted.

A diagram and additional information regarding the ACAIRN architecture running as an overlay on SuperNet is presented in Figure 7 and Table 2.

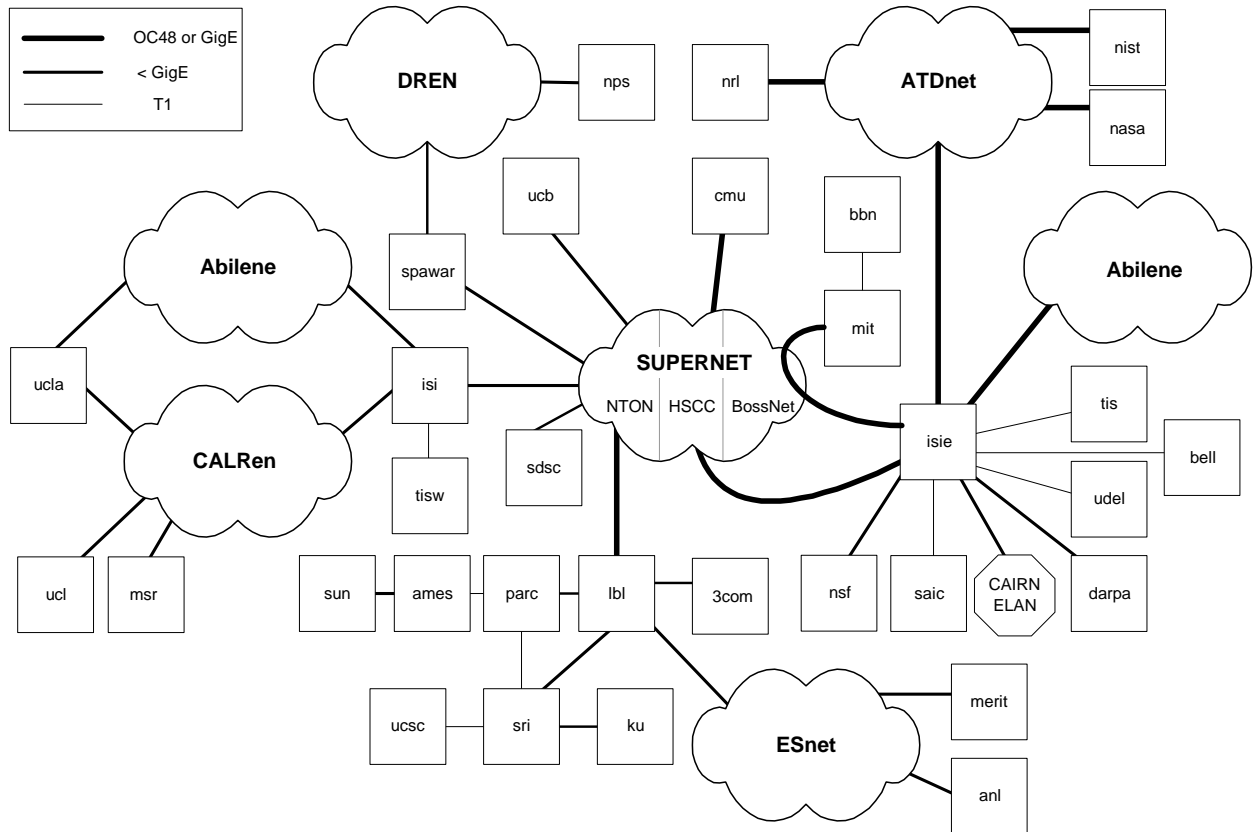


Figure 7 ACAIRN Infrastructure as an overlay on DARPA SuperNet

Table 2 ACAIRN Sites and Peering Networks

Acronym	Organization (Location)
Abilene	Internet2 Wide Area Network
ames	NASA AMES Research Center (California)
anl	Argonne National Labs (Chicago)
ATDNet	Advanced Technology Demonstration Network (Washington D.C. Metropolitan Area)
bbn	BBN (Massachusetts)
bell	Bell Laboratories (New Jersey)
CAIRN ELAN	CAIRN Emulated Local Area Network for ATM
CALRen	California Research Network
cmu	Carnegie Mellon University (Pennsylvania)
darpa	Defense Advanced Research Projects Agency (Virginia)
DREN	Defense Research and Engineering Network
ESnet	Energy Sciences Network POP (California)
isi	University of Southern California/Information Sciences Institute (California)
isie	University of Southern California/Information Sciences Institute East (Virginia)
ku	Kansas University (Kansas)
lbl	Lawrence Berkeley Laboratory (California)
merit	MERIT Networks
mit	Massachusetts Institute of Technology (Massachusetts)
msr	Microsoft Research

Acronym	Organization (Location)
nasa	National Aeronautics and Space Administration (Maryland)
netstar	Netstar
nist	National Institute of Standards and Technology (Maryland)
nps	Naval Postgraduate School (California)
nren	Naval Research and Education Network
nrl	Naval Research Laboratory (Maryland)
nsf	National Science Foundation
parc	Xerox Palo Alto Research Center (California)
saic	Science Applications International Corporation (Virginia)
sdsc	San Diego Supercomputing Center (California)
spawar	Space and Naval Warfare Systems Command (California)
sri	SRI International (California)
sun	Sun Microsystems (California)
tis	Trusted Information Systems (Maryland and California)
3Com	3Com Corporation
ucb	University of California at Berkeley (California)
ucl	University of College London (London)
ucla	University of California at Los Angeles (California)
ucsc	University of California at Santa Cruz (California)
udel	University of Delaware (Delaware)

Some of the research projects which utilized the ACAIRN network for research and experimentation activities are listed in Table 3.

Table 3 ACAIRN Research Activities

Research Project	Research Topics	Primary Organizations
Active Networks Backbone (ABone)	active network technology and testbed operations	USC/ISI,SRI,BBN, Trusted Information Systems (TIS)
Border Gateway Multicast Protocol (BGMP)	inter-domain multicast routing	Xerox Parc, USC/ISI, Microsoft Research
Bro: A System for Detecting Network Intruders in Real-Time	intrusion detection algorithms and technologies	Lawrence Berkeley National Laboratory (LBNL), The International Computer Science Institute (ICSI) Center for Internet Research (ICIR)
DARWIN	Developing protocols and algorithms that support service-oriented Internet (hierarchical packet scheduler, hierarchical CPU scheduler, tools for building overlay networks)	Carnegie Mellon University (CMU)
DNS Security (DNSSEC) in CAIRN	secure infrastructures for the Domain Name System	TIS, USC/ISI
Fault-Tolerant Mesh of Trust Applied to DNSSEC (FMESHD)	secure infrastructures for the Domain Name System	USC/ISI, Network Associates Laboratory (NAI Lab)

Research Project	Research Topics	Primary Organizations
Fault-Tolerant Networking Through Intrusion Identification and Secure Compartments (FNIISC)	scalable and robust fault-tolerant network applied to protecting the BGP routing infrastructure from faults and attacks	University of California at Los Angeles (UCLA), University of California at Davis (UC Davis), USC/ISI, NC State University (NCSU)
GLOBUS	Experimentation and deployment of GLOBUS GRID software and systems	ANL, USC/ISI
IPv6 Research and Experimentation	IPv6 deployment, testing, and peering	USC/ISI, Microsoft, 3-COM
Multicast Protocols Research and Experimentation	intra and inter domain multicast routing protocols, reliable multicast technologies, multicast dynamic traffic generators, multicast address allocation, PIM, IPv6 Multicast	USC/ISI, Xerox PARC, UCL, SAIC, SRI, UC Berkeley, UC Santa Cruz, UCLA
Multicast-based Inference of Network-internal Characteristics (MINC)	multicast-based estimators of the origins of degradation in network loss, delay and throughput	AT&T Center for Internet Research at ICSI (ACIRI), AT&T Labs-Research, University Of Massachusetts
National Internet Measurement Infrastructure (NIMI)	measurement technologies and systems to measure the global internet	National Laboratory for Applied Network Research (NLNR)
Network Time Synchronization Project (NTSP)	architecture, protocol and algorithms to synchronize computer clocks to national standard time	University of Delaware
Realizing Adaptive Distributive Internet Operations on ACTIVE Networks (RADIOACTIVE)	novel, hyper-scaleable applications and services using state-of-the-art application level and network level active networking technology	University College London (UCL)
Reliable Multicast Performance	development and test of protocols that use recovery hierarchies tree-based reliable multicast protocols	Carnegie Mellon University (CMU)
Quality of Service Research and Experimentation	RSVP implementation and testing, explicit congestion notification, CBQ, differentiated services, RED In-Out algorithms, DVMRP Flooding, ALT-Q	USC/ISI, LBL, UCLA, MIT, UCL, NASA, 3COM
Secure Border Gateway Protocol (S-BGP)	test and evaluation of technologies for securing BGP protocols	BBN Technologies
Secure Conferencing Access with Multicast Protocols for the Internet (SCAMPI)	novel mechanisms for supporting and utilizing Internet multicast conferencing	UCL

Research Project	Research Topics	Primary Organizations
SNMPv3	deployment and test of snmpv3	USC/ISI
Multiprotocol Label Switching (MPLS)	implementation of MPLS into CAIRN routers. MPLS experimentation and interoperability testing	NIST, USC/ISI
Video Teleconferencing	multicast VTC, MBONE peering	LBL, UCL, USC/ISI, SAIC, UCB, conference participation by majority of network users
X-Bone (Automated Overlay Network Deployment)	dynamically deployment and management of Internet overlays	USC/ISI

5.5.5. High Speed TCP over BossNet

The DARPA SuperNet network BossNet was utilized to conduct research into high speed end to end TCP. This research resulted in a successful demonstration of end to end TCP performance of 991 Mbps between Washington D.C. and Boston. This was an early demonstration of the ability of standard commercial off the shelf workstations and open source protocol stacks to perform at near gigabit per second rates across a wide area network. This demonstration required proper tuning of TCP protocol parameters and gigabit ethernet network interface card features.

5.5.6. Remote Media Immersion (RMI)

The USC Integrated Media Systems Center (IMSC) utilized SuperNet to demonstrate RMI technologies in a wide area environment. RMI is a breakthrough Internet technology for capturing, streaming, and rendering high-resolution digital video and multichannel audio. RMI transforms the Internet from a low-fidelity medium to a high-fidelity medium, delivering an experience beyond anything in existence today. The goal of the RMI is to reproduce the complete aural and visual ambience of an environment that includes people and other real and virtual elements. Immersive audio is a technique developed at IMSC for capturing the audio environment at a remote site, transmitting the information to a client location, and accurately reproducing (rendering) the complete audio sensation at the client location with full fidelity, dynamic range and directionality. Papers [7] and [8] describe some of the testing accomplished using SuperNet.

5.5.7. Meeting and Conference Support

Support for Principal Investigator, conferences and demonstrations were provided as part of the activities of the RESNETS II project. Some of these are described below.

5.5.7.1. Principal Investigator Meetings

Meetings and demonstrations: This included networking support for onsite demonstrations as well as high speed connections to the research network to allow individual researchers the ability to access their research facilities from remote meeting sites. The following meetings and conferences were supported.

2000 DARPA NGI PI Meeting

This meeting was held in McLean, Virginia on October 2-4, 2000. The meeting was held at a hotel and included a opportunity for the principal investigators to demonstrate their results. In order to support the demonstrations, the RESNETS II NOC provided advanced network services including high bandwidth connectivity to the SuperNet research testbed.

2001 DARPA NGI PI Meeting

This meeting was held in McLean, Virginia on January 5-7, 2002. The meeting was held at a hotel and included an opportunity for the principal investigators to demonstrate their results. In order to support the demonstrations, the RESNETS II NOC provided advanced network services including high bandwidth connectivity to the SuperNet research testbed. A diagram of the network configuration is shown in Figure 8.

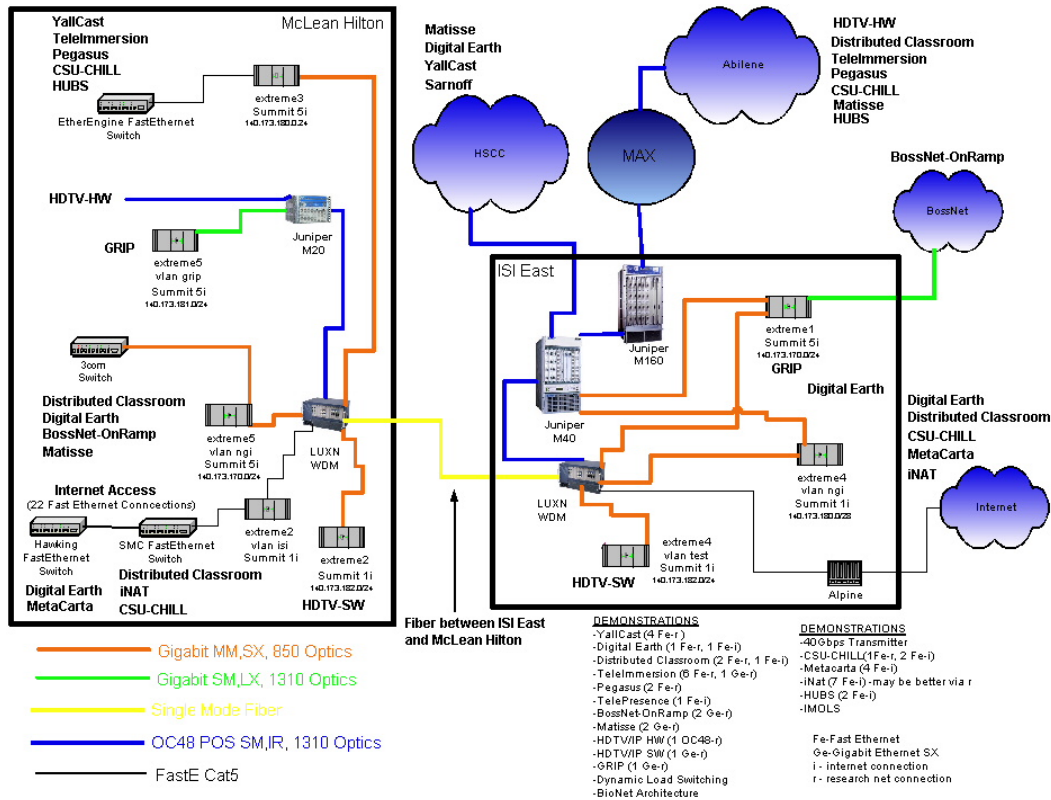


Figure 8 DARPA NGI Principal Investigators Meeting

5.5.7.2. Optical Fiber Conference

This Optical Fiber Conference was held in Baltimore Maryland in March 2000. SuperNet participated in OFC 2000 which included providing an at that time unprecedented amount of bandwidth to the show floor of the conference. Using the fiber architecture of BossNet, ATDnet and MONET, the hardware of a variety of vendors over 20 Gbps to the show floor, with over 17 Gbps of this capacity being utilized. Applications shown in the booth included a variety of HDTV over ATM applications, Network based MEMS device analysis, 980 Mbps single stream GbE TCP flow to Boston, Digital Earth and Internet analysis tools from CAIDA. A diagram of the network configuration is shown in Figure 9.

The network and applications presented during this conference were impressive enough to gather significant attention from the press which included a large picture and front page article on the business section of the Baltimore Sun newspaper.

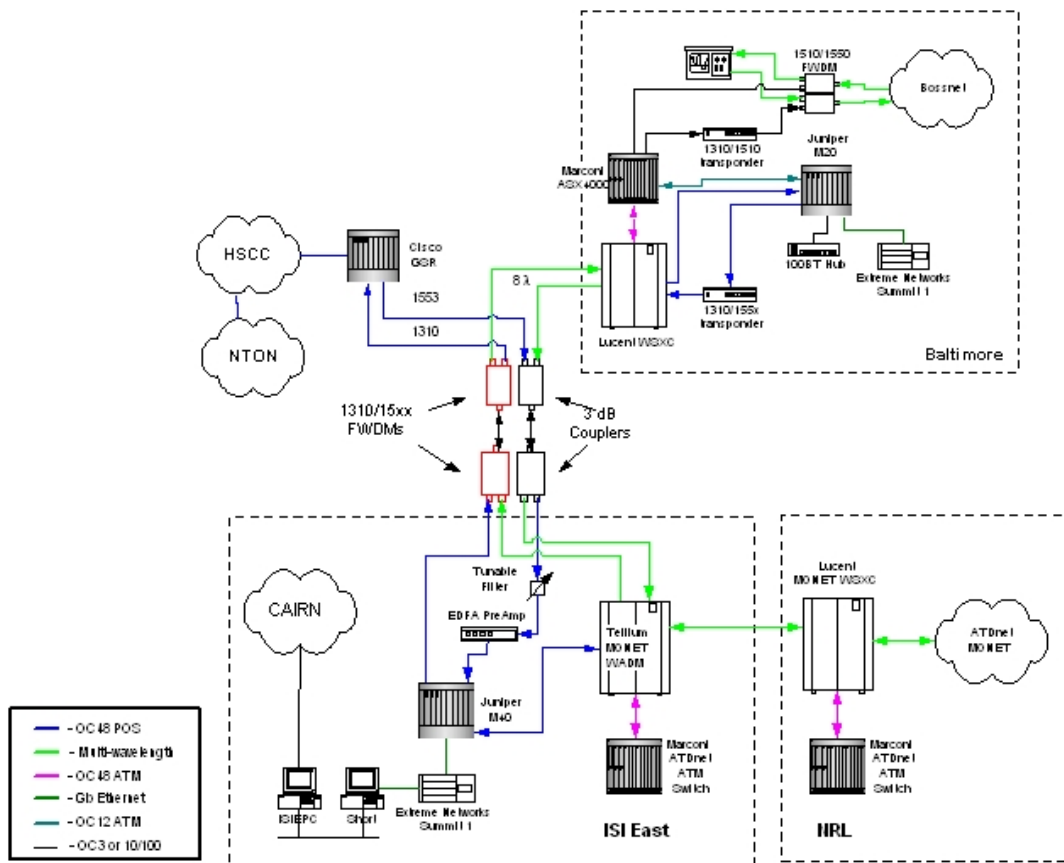


Figure 9 Optical Fiber Conference 2000 Network Configuration

5.5.7.3. SuperComputing Conferences

The annual SuperComputing Conferences are a venue of the Government and University Research and Education community to demonstrate and share results relating to high performance computing and networking. Activities supported via the RESNETSII program included active participation in Supercomputing 1999, 2000, 2001, and 2002. Activities included chairing of technical session on Network Protocol Tuning and conducting exhibits and demonstrations. As part of this conference are multiple exhibits where the latest technologies in high performance computing and networking are demonstrated. The DARPA SuperNet was an important part of these exhibits at the Supercomputing 2000 Conference held in Dallas, Texas in November 2000. This included providing key parts of the infrastructure for multiple research programs including the following:

- Accelerated Strategic Computing Initiative (ASCI)
- VisaPult:Image Based Rendering Assisted Volume Rendering - SC2000 Network Challenge Winner
- CalTech Particle Physics Using Globus
- Data Management Infrastructure for Climate Modeling Research (Striped FTP) - SC2000 Network Challenge Winner
- Stanford Linear Accelerator Center (SLAC)

- NASA Digital Sky Demo
- Digital Amphitheater
- Digital Earth

The SC2000 conference included a competition between high performance applications and network infrastructures to determine which could require and sustain the largest peak and aggregate bandwidth as part of their demo. This type of competition encourages researchers to stress the ability of their designs, computer equipment, network infrastructures and provides a good mechanism to foster technology improvements thru competition. As noted above, systems running over the DARPA SuperNet infrastructure won the Bandwidth Challenge event.

Bandwidth Challenge Winner at SC2000

The combination of Visapult and DARPA NGI SuperNet won the bandwidth challenge at SC2000 by demonstrating a peak bandwidth of 1.48 Gbit/s and a 60 minute average bandwidth of 582 Mbit/s. The Visapult application allowed for the remote visualization of an 80 GByte data set. The data set was located at LBL in Berkeley California. The real time visualization was rendered by an 8 CPU SGI Origin on the showroom floor in Dallas, Texas.

Similar research experiments were conducted during Supercomputing 1999 and 2001.

Additionally, the RESNETSII project provided all networking for the National Coordination Office (NCO) for Supercomputing 2001. This included support for NCO sponsored research from NASA, NIH, Tektronix, and USC/ISI. This included over 5 Gbits/s to the showroom floor booth. Demonstrations include Uncompressed HDTV/IP, and Gigabit Rate IP Security. The Uncompressed HDTV/IP application generated a 1.6 Gbit/s stream being sent from Seattle to Denver Colorado. A diagram of the network and NCO booth schematic is at SC2001 is show in Figure 10.

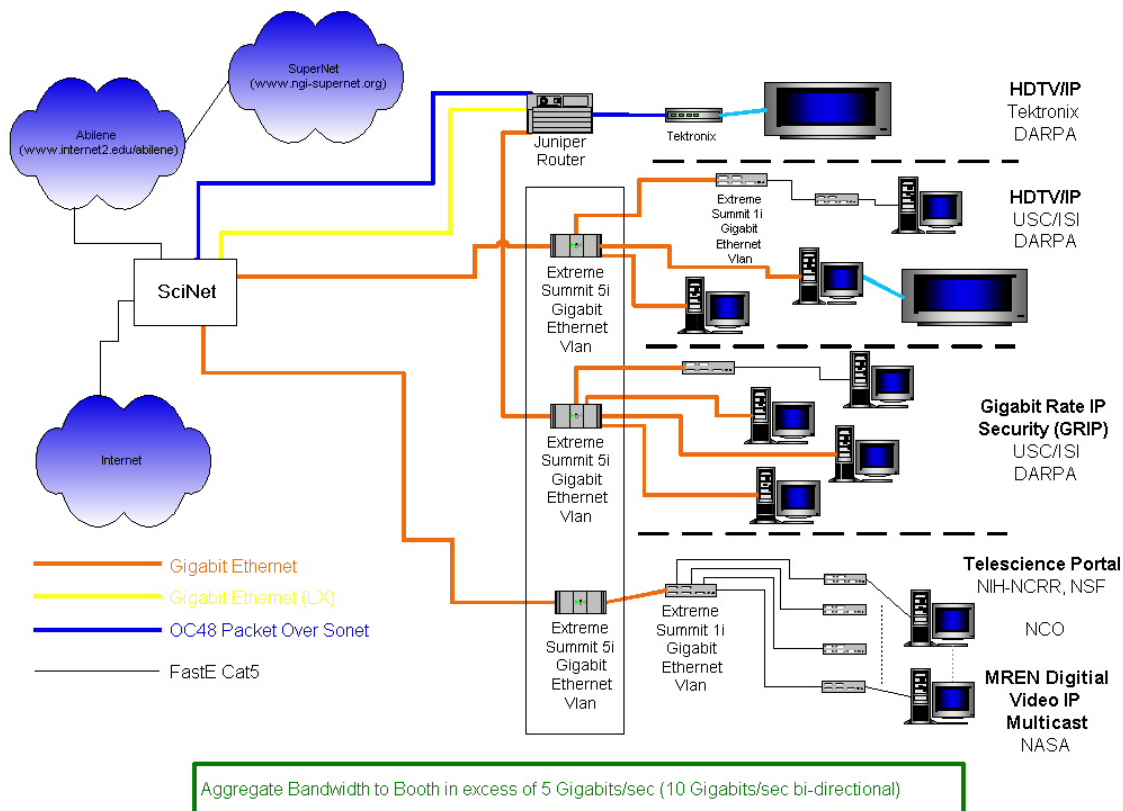


Figure 10 SuperComputing 2001 National Coordination Office (NCO) Booth

5.5.8. Other Meetings and Conferences Supported

There were many other meetings and conferences supported as part of the RESNETSII project. Some were ongoing regularly scheduled, and others were one time events. These included the following:

- Mid-Atlantic Crossroads (MAX) Engineering and Management Meetings
- ATDNet Steering Committee Meetings
- ATDNet Experiments Meetings
 - Participated in ATDNet experimental activities and meetings
- NTON User meetings
- VLBI meetings
- WNS meetings
- JET meeting
 - Joint Engineering Taskforce (JET) meetings are held monthly at NSF. The JET meetings bring together multiple DOD, Government, and Education research networks for technical and programmatic exchange.
- DARPA,LSN, and NSF Distinguished Lecture Series
 - Provided network support and real media archive for this lecture series
- AMES Gigabit Workshop
 - Presented at the AMES Gigabit Workshop in August 2000 on the subject of End-to-End performance across High Bandwidth Delay Product Networks
- Internet2 Multicast Workshop
 - Hosted the Internet2 Multicast Workshop at ISI East

5.6. Advanced Network and Application Research

RESNETSII project activities also included direct conduct of and participation in advanced network and high performance application research.

5.6.1. End-to-End Performance across High Bandwidth Delay-Product Networks

A key objective for the RESNETSII project was the integration of multiple and diverse regional networks into an infrastructure where researchers could collaborate as well as conduct research experiments not possible when confined to just their regional network. The combination of multiple regional infrastructures into a interconnected wide area national infrastructure enabled research collaborations and experiments not possible across the individual regional networks. One example of this is the research conducted into application network performance across large bandwidth delay product networks. In the timeframe of the RESNETS project, large bandwidth was considered to be 1 Gigabit/s or larger. Large delay networks were associated with packet round trip times of greater than 70 ms which is typical for high performance cross country networks. Conducting performance testing in this environment demonstrated many deficiencies within network protocols and application designs which were not evident from testing across regional infrastructures.

There are several issues which result in poor performance of TCP streams when operated in this environment. These issues include: i) TCP congestion control response which is inappropriate and incompatible with the loss and reorder rates observed for these network conditions and ii) the advent of large bandwidth UDP streams which do not have congestion control and do not share the network resources fairly. The goal of this research was to characterize these issues, develop solutions, and verify via experiment.

We developed systems, protocol tuning procedures, and experiment procedures to push the limits of end-to-end performance in this environment. A summary of these techniques and activities are described below.

The purpose of the network performance testing was to determine the maximum end system to end system performance possible across the network path to be used for a specific experiment or demonstration. This testing demonstrated the many issues associated with the conduct of high speed end to end data communications across high bandwidth-delay product networks. The combination of high bandwidth and long round trip times can be problematic for reliable transport protocols such as TCP. These issues revolve around hosts and TCP congestion control responses to packet loss and reordering. The parameters across the WAN are significantly different compared to that observed in lab testing. The result is that applications that perform well in the lab sometimes have less success when operated across the wide area.

Many of the end systems available included gigabit ethernet network interface cards. Therefore it was desired was to realize TCP flows on this order even for paths with long round trip times. With the proper host tuning, this was achieved. The host and network tuning required to achieve these results are summarized below:

- Path MTU Discovery. Host systems must utilize Path MTU Discovery per RFC1191. This allows hosts to discover the optimum MTU size to use for the current end to end path. For the SuperNet testing MTU sizes of 4470 and 9000 bytes were available.
- Large Window Extensions. TCP includes definition of a window parameter which defines how much data may be in the network pipe prior to receipt of specific acknowledgement. The bandwidth-delay product of the network path determines the optimum value of this parameter. For the SuperNet testing this value was high enough that host system support of RFC1323 "Large Windows" extensions to TCP was required.
- Large Socket Buffers. The host system must support large enough socket buffers for reading and writing data to the network. These values needed to be in line with the bandwidth-delay product. For the SuperNet testing a conservative value of 3-8 MB (megabytes) was used. Typical host default values are 64 kB (kilobytes) .
- TCP Selective Acknowledgments (SACK). TCP SACK (RFC2018) allows for greater tolerance to packet reordering. Data communications across the wide area network results in increased packet reorder events so this feature is more important as compared to local area communications.

Testing of network performance was conducted using iperf[9]. This application provides a basic test of the end to end performance. The data transfer is from host memory to host memory, so minimum host resources (disk drives, memory copies, CPU processing) are exercised. Results achieved with iperf are generally best case, since real applications will require disk access, memory copies, and additional application CPU utilization. Below is a typical iperf output from a network performance test.

Performance Test Result: MIT/LL (Boston,MA) --->ISI East (Arlington, VA)
[root@kame root]# iperf -s -w 8M -i 5

Server listening on TCP port 5001

TCP window size: 16.0 MByte (WARNING: requested 8.0 MByte)

[6] local 140.173.180.5 port 5001 connected with 140.173.174.10 port 32822
[ID] Interval Transfer Bandwidth
[6] 0.0- 5.0 sec 577 MBytes 922 Mbits/sec
[6] 5.0-10.0 sec 584 MBytes 935 Mbits/sec
[6] 10.0-15.0 sec 584 MBytes 935 Mbits/sec
[6] 15.0-20.0 sec 584 MBytes 935 Mbits/sec
[6] 20.0-25.0 sec 584 MBytes 935 Mbits/sec
[6] 25.0-30.0 sec 584 MBytes 934 Mbits/sec
[6] 0.0-30.1 sec 3.4 GBytes 933 Mbits/sec

An example of our results in this area was shown as a result of winning the Internet2 Land Speed Record competition which is described below.

Internet2 Land Speed Record

The Internet2 Land Speed Record (I2-LSR) competition for the highest-bandwidth, end-to-end networks is an open and ongoing contest. Internet2 Land Speed Record entries are judged on a combination of how much bandwidth they used and how much distance they covered end-to-end, using standard Internet (TCP/IP) protocols.

The first Internet2 Land Speed Record competition results were announced in March 2000 and the winning team used the DARPA SuperNet. A new standard for Internet performance was established by transferring 8.4 GB worth of data from Redmond, Washington to Arlington, Virginia (5,626 Km) in 81 seconds at a rate of over 830 megabits per second. This demonstration won both the single stream and multistream classes of the 2000 Internet2 Land Speed Record competition. Details of history of this competition are available from the Internet2 website[10].

5.6.2. Packet Loss, Reordering, and Impact on Transport Protocols

We conducted research into packet loss and reordering across commercial IP backbones. This type of information is very important to end system performance and can greatly impact protocol performance such as the Transmission Control Protocol (TCP). We constructed a test setup using the SuperNet infrastructure to allow us to inject high bandwidth flows into the network and observe the amount of packet loss and reordering across the wide area infrastructure. The experiment setup is shown in Figure 11.

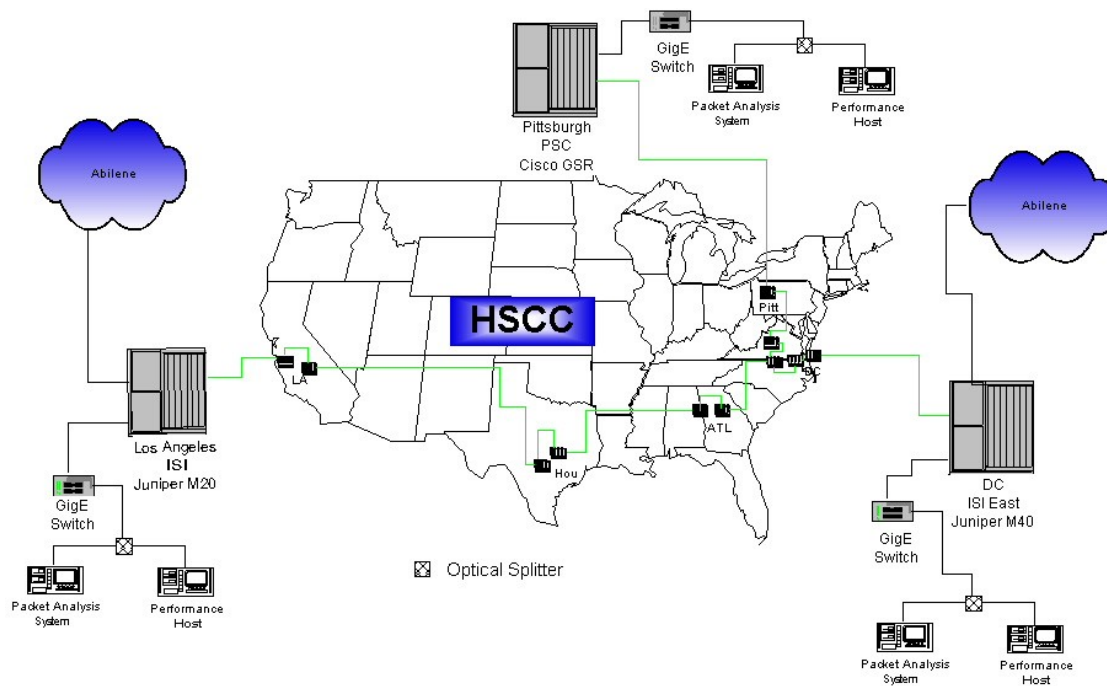


Figure 11 Packet Loss and Reordering Research Infrastructure

We conducted end-to-end measurements of UDP flows and characterized the loss and reordering behavior. In total we collected data on over 350 UDP flows, comprising approximately 120 million packets, over the three different network paths, at various MTUs and bandwidths (up to 900 Mbps).

Our observations indicate that packet loss in the backbone is a rare event. Packet reordering is relatively common at the flow level: almost half of the tests saw some instance of packet reordering. Despite this, the

fraction of packets reordered is low, with the overwhelming majority of flows seeing less than 1 packet in 1000 reordered.

We evaluated current methodologies for classification of reordering, and demonstrate the importance of choosing an appropriate metric for the application or protocol. For example, the data shows that a flow with less than 0.4% reordering (measured by a simple percent metric) can generate either no spurious congestion events for TCP, or several thousand spurious congestion events, depending on the pattern of reordering. The resulting impact to TCP performance can be large.

In addition, the effects of congestion control on high bandwidth flows were studied using the SuperNet infrastructure. Congestion control algorithms are implemented by TCP based applications as a vehicle to allow all users to share the network fairly. Since in a packet switched network the routers and links are shared amongst multiple users, a mechanism to ensure an acceptable use profile is needed. However, the majority of User Datagram Protocol (UDP) applications do not utilize a congestion control mechanism.

Since the majority of internet traffic is TCP based, this has not proved to be a problem to date. However, with the advent of high bandwidth multimedia applications, this is a problem that will need to eventually be addressed. We utilized the SuperNet infrastructure to conduct experiments on such UDP applications. The

measurements demonstrated that congestion control schemes utilized in TCP are not necessarily appropriate for multimedia and other UDP based applications. Details of this testing and analysis were presented in a paper "Implementing Congestion Control in the Real World" [11].

5.6.3. HDTV over IP

The SuperNet infrastructure was utilized to develop technologies to allow uncompressed HDTV to be distributed across standard IP networks. Transmitting of the 1.6 Gbit/s HDTV data stream cross the wide area network demonstrated the feasibility of using IP networks for high bandwidth data.

A system was designed to transports uncompressed HDTV content over IP networks. This system was motivated by growth in use of digital video, and the ever increasing capacity of local- and wide-area Internet links. This system was used to demonstrate the feasibility of IP as a transport for very high quality video, and to highlight areas where performance bottlenecks exist, and where further development is needed. The system was constructed from commodity components, and was tested over existing commercial IP backbone networks. Performance was shown to be good, with the end system being the main limiting factor. The results of this system development and associated testing over SuperNet are reported in [12].

This work also contributed to work on two IETF Request for Comments (RFCs) related to UDP and multimedia applications across IP networks, reference [13] and [14].

5.6.4. Extended Gigabit Ethernet WAN Services

Experimentation on SuperNet was used to transported Gigabit Ethernet packets directly over WDM for 1062 km in the successfully deployed MONET Washington DC network. The experiment was an early indicated that the Gigabit Ethernet can be used for WAN or backbone transmission bypassing the ATM and SONET intermediate protocol layers. It also indicated that the physical layer of the Fibre Channel, which is adopted by the Gigabit Ethernet, can be used as a long-haul transmission protocol. Figure 12 shows the resulting optical eye diagrams after different transmission distances.

The detailed results were documented in the paper Transport of Gigabit Ethernet Directly over WDM for 1062 km in the Monet Washington DC Network [15].

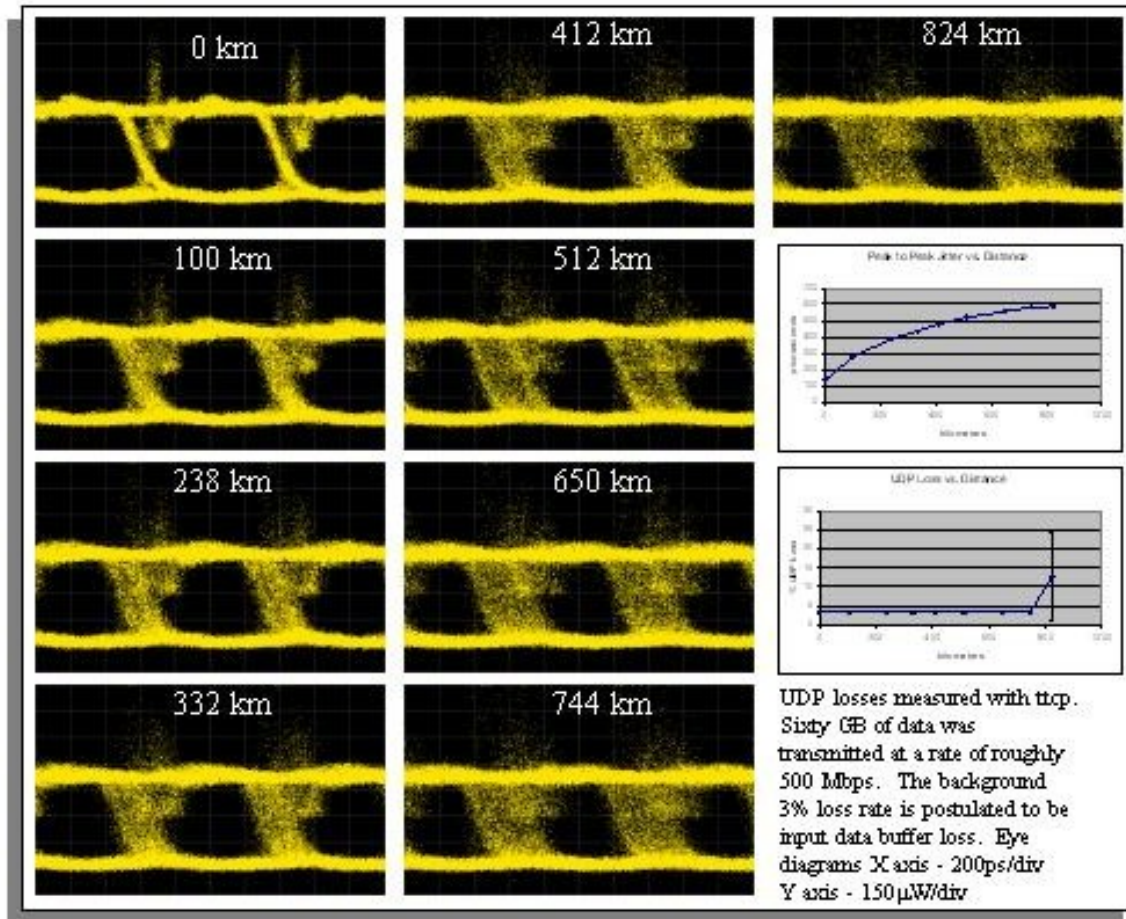


Figure 12 Eye Diagrams of Gigabit Ethernet Channel at Multiple Transmission Distances

5.6.5. High Speed Optical Transmission Research

The BoSSNet and ATDNet regional network infrastructures were optical WDM networks which allowed for research and investigation into issues with all optical data transmissions. This type of data transmission is distinct from common data networks which may use optical segments, but generally regenerate the data packet via an optical-electrical-optical conversion process at intermediate network elements. While use of all optical technology has great promise in terms of increased performance, greater flexibility, and reduced cost there are many technical issues which require further research and development. An important research area is characterization and compensation for the effects of linear and non-linear transmission impairments associated with optical data transmission.

Equipment was placed at the DARPA node of ATDNet to allow optically transparent (i.e., no optical-electrical-optical conversions) on multiple wavelengths between any ATDNet site and MIT Lincoln Laboratory in Boston, Massachusetts. The purpose of this infrastructure was to; i) test novel concepts for high speed digital and analog data transmission, ii) evaluate technical issues associated with long distance all optical transmission, iii) provide data for initial thinking on how these issues might impact new concepts for data transmission and rapid provisioning (such as intelligent systems based on optical burst switching and/or lambda switching). Research was conducted to realize channel rates of 10 to 40 Gigabits/sec. The effort includes the optimization of performance in the presence of optical non-linear, polarization, and transient effects. This work was a collaboration between the Laboratory for Telecommunication Sciences (LTS) at NSA, Telcordia, UMBC, and ISI.

As part of the RESNETS II project, we participated in the conduct of experiments across the BoSSNet and ATDNet infrastructures to better understand these phenomena. Some of these results were reported in a paper titled "Timing Jitter Due to Intra-Channel Nonlinearities in an Installed Fiber Optic Network [16]. This paper presented a comparison between numerical simulation and field measurements of intra-channel nonlinearity-induced timing jitter. Both simulation and measured data demonstrate that the timing jitter is strongly dependent on the distribution of dispersion compensation.

6. Important Findings and Conclusions

The instantiation of the research networks described in this report continued a successful history of DARPA supported network research testbeds. Prior DARPA supported testbeds included ARPANet, DARTNet, CAIRN, and ACAIRN (which eventually ran as an overlay on DARPA SuperNet).

Like its predecessors, the SuperNet testbed demonstrated that network research testbeds provide a valuable vehicle for important research topics to be deployed, tested, and refined. Additionally, this process provides a method for moving research topics out of the laboratories and into government and commercial deployment. As an example of this, several of the infrastructure components which were part of the SuperNet and DARPA NGI program have transitioned to other government agencies where they will play a key role in the development and deployment of advanced network technologies. This includes the incorporation of BoSSNET and ATDNet into the Global Information Grid (GIG) Testbed which is supported by the Transformational Communications Office (TCO) as part of the Transformational Communications Architecture (TCA) program. Incorporation of these DARPA developed resources into these programs is a testament to the value of the DARPA supported research programs.

The activities of the RESNETSII project included integrating the regional testbeds, assisting researchers, providing network operations, and conducting research. These activities provided a good position from which to comment on the overall impact of these types of research community testbeds.

As already mentioned, these types of testbeds provide a valuable vehicle for important research topics to be deployed, tested, and refined. However, successful transition of a technology from a testbed environment to wider acceptance and use is a difficult task. Key factors which influence the success of the transition are the timing of the research topic in relation to current commercial practices, engagement of a vibrant community of researchers, and access to appropriate resources for thorough research experimentation.

An important characteristic of a network research testbed is that the research agenda be driven by the research coming from the networking community. Toward this end the architecture of the network needs to facilitate enough flexibility to respond to new and changing research agendas. This often requires open hardware and software systems, flexible infrastructures, and ability to rapidly reprogram and reconfigure systems.

Network operations, management, and monitoring can represent significant costs of the overall testbed. This should be factored in from the very early planning stages. Networks are increasingly complicated and the skill level and people resources are increasing accordingly.

The network research community has been very prolific at developing many new technologies. Many technologies have migrated to production networks and many are still in some stage of research and/or development. The resulting complexity in today's network infrastructures demands an increased attention to testbed management tools and instrumentation. This is one area where best practices of several years ago will not be sufficient as one looks to the future of testbeds or production networks. Increased investment in testbed management tools and instrumentation will be required for future testbed efforts.

The social issues can dominate the probability of success or failure for a distributed testbed effort. Careful attention and nurturing of a vibrant research community is the best way to maximize successful use of a testbed resource.

7. Implications for Further Research

Network testbeds have proven that they can play a critical in the advancement of research agendas and development of technologies. Key factors to maximize probability of success are:

- identification of a technology area that is well positioned to benefit from testbed activities
- identification and support of a research community where the benefits of collaboration, available testbed resources, and technology focus areas all result in a synergistic environment
- provision of adequate resources for management, monitoring, and instrumentation of network infrastructure

Lessons learned from the SuperNet testbed indicate that as network technologies (and subsequent testbeds) become more complicated, increasing emphasis and resources will need to be applied to advanced management tools and instrumentation. This will be required to provide the researchers with the flexibility and information required to utilize the testbed resources. This will likely require long-term funding to maximize the time a mature well functioning testbed is available to the research community.

With the advent of new technologies there will be increasing options for utilizing existing infrastructures to partition dedicated network resources for the purpose of creating network research testbeds. In particular, the maturing of optical network technologies will allow increasingly creative, rich, and complicated testbed topologies to be created at significantly less cost than previous testbeds efforts which required purchase of truly dedicated network resources. These new brand of “overlay networks” have the potential to provide many opportunities for research to faster advance from simulation, to prototype, to deployment.

However, significant advancements in the area of Network Management, Resource Provisioning, and Network Control will be required to fully realize the potential of these advanced network technologies.

A strong recommendation is made for further investment in the areas of intelligent Network Management, Resource Provisioning, and Network Control.

8. References

- 1 SuperNet Research Network, <http://www.ngi-supernet.org>
- 2 Netarchive Data Plotter, <http://www.ittc.ku.edu/aai/plot/Main.html>
- 3 Multi Router Traffic Grapher (MRTG), <http://mrtg.hdl.com/mrtg.html>
- 4 K.I. Kellerman and A.R. Thompson. The Very-Long-Baseline Array, Scientific American, Vol.258 No.1 pages. 54-63,1988.4
- 5 Whitney, A.R., et al, "The Gbps e-VLBI Demonstration Project". ftp://web.haystack.edu/pub/e-vlbi/demo_report.pdf.
- 6 Christian Rembe, Rishi Kant, Michael P. Young, Richard S. Muller, "Network-connected MEMS-measuring system for high-speed data transfer to CAD and simulation tools," Conference on Vibration Measurements by Laser Techniques," Italian Assn. for Laser Velocimetry, Ancona, Italy, 18-21 June 2002
- 7 Yima: A Second Generation Continuous Media Server. Cyrus Shahabi, Roger Zimmermann, Kun Fu, and Shu-Yuen Didi Yao. Published in the IEEE Computer magazine, June 2002, pp. 56-64. <http://idefix.usc.edu/pubs/IEEEComp.pdf>
- 8 Retransmission-Based Error Control in a Many-to-Many Client-Server Environment. Roger Zimmermann, Kun Fu, Nitin Nahata, and Cyrus Shahabi. Accepted for presentation at the SPIE Conference on Multimedia Computing and Networking 2003(MMCN 2003), Santa Clara, California, January 29-31, 2003.
- 9 National Laboratory for Applied Network Research: Network performance measuring tool, <http://dast.nlanr.net/Projects/Iperf/> (2003)
- 10 Internet2 Land Speed Record History, <http://lsr.internet2.edu/history.html>
- 11 Ladan Gharai & Colin Perkins, Implementing Congestion Control in the Real World, Proceedings of the IEEE International Conference on Multimedia and Expo, Lausanne, Switzerland, August 2002.
- 12 Colin Perkins, Ladan Gharai, Tom Lehman & Allison Mankin, Experiments with Delivery of HDTV over IP Networks, Proceedings of the 12th International Packet Video Workshop, Pittsburgh, April 2002.
- 13 Ladan Gharai, Colin Perkins, Gary Goncher and Allison Mankin, RTP Payload Format for Society of Motion Picture and Television Engineers (SMPTE) 292M Video , Internet Engineering Task Force, RFC3497, March 2003.
- 14 Ladan Gharai and Colin Perkins, RTP Payload Format for Uncompressed Video, Internet Engineering Task Force, in Last Call, January 2004
- 15 W. Xin, G.K. Chang, and T.T. Gibbons, "Transport of gigabit Ethernet directly over WDM for 1062 km in the MONET Washington DC network", paper WC1.4, Proc. IEEE summer topical meeting, Boca Raton, FL (2000)
- 16 Lenihan,Tudury,Carter,Sinkin,Marks,Menyuk,Runser,Goldman,Lehman, "Timing Jitter Due to Intra-Channel Nonlinearities on an Installed Fiber Optic Network", Conference on Lasers and Electro Optics/International Quantum Electronics Conference (CLEO/IQEO) 2004.